

REMARKS

Upon entry of this Amendment, claims 1-21 will remain pending. By this amendment, Applicant amends the specification to improve form or grammar.

In the Office Action,¹ the Examiner rejected claims 1-4, 7-9, and 13-20 under 35 U.S.C. § 102(b) as being anticipated by “Two-level Processor Scheduling for Multiprogrammed NUMA Multiprocessors” to Fukuda et al. (“Fukuda”); and rejected claims 5, 6, 10-12, and 21 under 35 U.S.C. § 103(a) as being unpatentable over Fukuda in view of U.S. Patent No. 5,881,284 to Kubo (“Kubo”).

Applicants thank the Examiner for withdrawing the previous rejections under 35 U.S.C. §§ 102(b), 103(a), 112, and 101. Applicants respectfully traverse the Examiner’s rejections for the following reasons.

I. The rejection of claims 1-4, 7-9, and 13-20 under 35 U.S.C. § 102(b) as being anticipated by Fukuda is improper.

Applicants respectfully traverse the rejection of claims 1-4, 7-9, and 13-20 under 35 U.S.C. § 102(b) as being anticipated by Fukuda. In order to properly establish that Fukuda anticipates Applicants’ claimed invention under 35 U.S.C. § 102, each and every element of each of the claims in issue must be found, either expressly described or under principles of inherency, in that single reference. Furthermore, “[t]he identical invention must be shown in as complete detail as is contained in the ... claim.” See M.P.E.P. § 2131, quoting *Richardson v. Suzuki Motor Co.*, 868 F.2d 1126, 1236, 9 U.S.P.Q.2d 1913, 1920 (Fed. Cir. 1989).

¹ The Office Action contains a number of statements reflecting characterizations of the related art and the claims. Regardless of whether any such statement is identified herein, Applicants decline to automatically subscribe to any statement or characterization in the Office Action.

Fukuda does not disclose each and every element of Applicants' claimed invention. Claim 1 calls for a combination including, for example, "a topology monitoring unit for monitoring a status of the CPUs and generating status information signals indicative of the status of each group of node boards" (emphasis added). Fukuda fails to teach or suggest at least these elements.

The Examiner asserts that Fukuda's group server constitutes the claimed "topology monitoring unit." Office Action at p. 3. Applicants disagree. Fukuda's group server is "for creating a processor group and allocating the group to a new process." Fukuda at p. 344, section 2.2. Fukuda discloses: "when a private scheduler requires more processors, it issues this request to the group server." Fukuda at p. 345 (emphasis added). Receiving by the group server a request for allocation of processors, as taught by Fukuda, does not constitute a teaching or suggestion of "a topology monitoring unit for monitoring the status of the CPUs," as recited by claim 1.

Moreover, Fukuda's group server does not generate "status information signals indicative of the status of each group of node boards," as recited by claim 1. Accordingly, Fukuda's group server cannot constitute the claimed "topology monitoring unit."

Claim 1 also recites a combination including, for example, "a job scheduling unit for receiving said status information signals and a job, and scheduling the job to one group of node boards on the basis of which group of node boards have the resources required to execute the job as indicated by the status information signals" (emphasis added). The Examiner has not indicated how Fukuda allegedly teaches or suggests the claimed "job scheduling unit." Fukuda discloses: "a high-level scheduler (a global

scheduler) and multiple low-level schedulers (private schedulers).” Fukuda at p. 344.

Neither Fukuda’s high-level scheduler nor Fukuda’s low-level schedulers constitutes the claimed “job scheduling unit.”

Regarding Fukuda’s high-level scheduler (group server), as discussed above, Fukuda’s group server does not receive “status information signals.” Further, Fukuda’s group server does not schedule “the job to one group of node boards on the basis of which group of node boards have the resources required to execute the job as indicated by the status information signals,” as recited by claim 1. Fukuda’s group server simply receives and fulfills a request from the private (low-level) scheduler to allocate processors to a processor group. Fukuda at p. 345, section 2.2(1), paragraph 3. The request itself indicates “the number of processors to be allocated” as determined by the private (low-level) scheduler. A group server receiving a request that indicates a number of processors to allocate, as taught by Fukuda, does not constitute a teaching or suggestion of “receiving said status information signals and a job, and scheduling the job to one group of node boards on the basis of which group of node boards have the resources required to execute the job as indicated by the status information signals,” as recited by claim 1 (emphasis added). Accordingly, Fukuda’s group server cannot constitute the claimed “job scheduling unit.”

Fukuda’s low-level (private) scheduler also cannot constitute the claimed “job scheduling unit.” Fukuda discloses: “each private scheduler is associated with a processor group.” Fukuda at p. 344, section 2.1(2). Fukuda’s private schedulers therefore monitor only one group, the group that the private scheduler is associated with. Monitoring a single group, as disclosed by Fukuda’s private schedulers, does not

constitute a teaching or suggestion of receiving “status information signals indicative of the status of each group of node boards.”

In addition, Fukuda's low-level schedulers do not “schedule a job,” as recited by claim 1. Rather, Fukuda's group server creates “a processor group and allocat[es] the group to a new process [job].” Fukuda at p. 344, section 2.2(1); p. 343, section 1, paragraph 1.

Moreover, Fukuda's low-level scheduler does not “schedule the job to one group of node boards on the basis of which group of node boards have the resources required to execute the job as indicated by the status information signals,” as recited by claim 1. Rather, Fukuda's low-level scheduler determines when more processors are required within the group that the scheduler is responsible for and “issues a request [for the required processors] to the group server.” Accordingly, Fukuda's low-level schedulers also cannot constitute the claimed “job scheduling unit.”

Because Fukuda does not teach or suggest each and every element recited by claim 1, Fukuda cannot anticipate this claim. Independent claim 13, although of different scope, recites elements similar to those recited by claim 1. Claims 2-4, 7-9, and 14-20 depend from claims 1 and 13 and therefore include all of the elements recited therein. Accordingly, for at least the reasons discussed above with respect to claim 1, Fukuda cannot anticipate claims 2-4, 7-10, and 13-20. Applicants therefore respectfully request the Examiner to reconsider and withdraw the rejection of claims 1-4, 7-9, and 13-20 under 35 U.S.C. § 102(b) as being anticipated by Fukuda.

Dependent claim 3 recites a combination including, for example, “the status information signals for each group of node boards indicate a number of CPUs available

at each radius; and wherein the job requires resources having a required radius to execute the job, and the job scheduling unit allocates the job to the one group of node boards on the basis of which group of node boards have CPUs available with a radius corresponding to the required radius" (emphasis added). The Examiner asserts that Fukuda teaches these elements on pages 344-345, sub-section 2.2. Office Action at p. 3. This assertion is incorrect. Fukuda does not teach or suggest "status information signals for each group of node boards [that] indicate a number of CPUs available at each radius," as recited by claim 3. Moreover, neither the cited portion of Fukuda nor any other portion teaches or suggests a job that "requires resources having a required radius to execute the job," or allocation of a job "on the basis of which group of node boards have CPUs available with a radius corresponding to the required radius." Accordingly, for at least this additional reason, Fukuda cannot anticipate dependent claim 3.

II. The rejection of claims 5, 6, 10-12, and 21 under 35 U.S.C. § 103(a) as being unpatentable over Fukuda in view of Kuba is improper.

Applicants respectfully request that the Examiner reconsider and withdraw the rejection of claims 5, 6, 10-12, and 21 because a *prima facie* case of obviousness has not been established with respect to these claims.

To establish a *prima facie* case of obviousness, three basic criteria must be met. First, there must be some suggestion or motivation, either in the references themselves or in the knowledge generally available to one of ordinary skill in the art, to modify the reference or to combine reference teachings. Second, there must be a reasonable expectation of success. Finally, the prior art reference (or references when combined) must teach or suggest all the claim limitations. The teaching or suggestion to make the claimed combination and the reasonable expectation of success must both be found in the prior art, and not based on applicant's disclosure. *In*

re Vaeck, 947 F.2d 488, 20 USPQ2d 1438 (Fed. Cir. 1991). M.P.E.P.
§ 2142, 8th Ed., Rev. 2 (May 2004), p. 2100-128.

A prima facie case of obviousness has not been established because, among other things, neither Fukuda nor Kuba teaches or suggests each and every element recited by Applicants' claims.

Claims 5 and 6 depend from claim 1 and claim 21 depends from independent claim 13. As discussed above, Fukuda fails to teach or suggest each and every element recited by independent claims 1 and 13 and required by dependent claims 5, 6, and 21. Kubo fails to cure the deficiencies of Fukuda discussed above, nor does the Examiner rely on Kubo for such teachings. Kubo does not teach or suggest "a topology monitoring unit for monitoring a status of the CPUs and generating status information signals indicative of the status of each group of node boards," or "a job scheduling unit for receiving said status information signals and a job, and scheduling the job to one group of node boards on the basis of which group of node boards have the resources required to execute the job as indicated by the status information signals," as recited by claim 1 and required by dependent claims 5 and 6 (emphasis added).

Accordingly, the cited references fail to teach each and every element recited in claims 1 and 13 and required by claims 5, 6, and 21. Thus, for at least the reasons discussed above with respect to claims 1 and 13, no prima facie case of obviousness has been established for claims 5, 6, and 21. Applicants therefore respectfully request that the Examiner reconsider and withdraw the rejection of claims 5, 6, and 21 under 35 U.S.C. § 103(a) as being unpatentable over Fukuda in view of Kubo.

Further, dependent claim 5 recites a combination including, for example, “status information signals [that] include, for each host, a number of CPUs which are available for each radius; and ... [a] scheduling unit [that] maps the job to a selected host having a maximum number of CPUs available at a radius corresponding to the required radius for the job” (emphasis added). As discussed above, Fukuda does not teach or suggest determination of “a number of CPUs which are available for each radius.” Moreover, the Examiner concedes “Fukuda did not specifically teach[] the scheduling unit maps the job to a selected host having a maximum number of CPUs available at a radius corresponding to the required radius for the job.” Office Action at p. 5. The Examiner nevertheless asserts that Kubo discloses this element in col. 3, lines 8-52. Id. at p. 6. This assertion is incorrect.

The cited portion of Kubo discloses: “measurement mechanism 1 for measuring a utilization of a cluster.” Kubo, col. 3: 10. Kubo’s “resource utilization is, for example, the use ratio of processors, the use ratio of channels, the ratio of an area used in the main storage, a paging frequency, and input/output operation frequency, [and] the use ratio of a software resource.” Kubo, col. 3:26-31. None of these utilization measurements disclosed by Kubo constitutes a teaching or suggestion of scheduling “on the basis of which group of node boards have CPUs available with a radius corresponding to the required radius,” as recited by claim 5.

Because Fukuda and Kubo, taken alone or in combination, fail to teach or suggest each and every element recited by claim 5, no prima facie case of obviousness has been established with respect to this claim. Applicants therefore respectfully

request that the Examiner reconsider and withdraw the rejection of claim 5 under 35 U.S.C. § 103(a) as being unpatentable over Fukuda in view of Kubo.

Independent claim 10, although of different scope, recites elements similar to those recited by claim 1. Fukuda therefore fails to teach or suggest each and every element recited by claim 10 for at least the reasons discussed above. Further, independent claim 10 recites a combination including, for example,

(d) scheduling the job to the resources which are available to execute the job as based on the status information signals and the physical topology, and the resources required to execute the job.

Fukuda fails to teach or suggest “scheduling ... based on ... the physical topology,” as recited by claim 10. Kubo fails to cure these deficiencies of Fukuda, nor does the Examiner rely on Kubo for such teachings. The Examiner relies on Kubo for allegedly disclosing “comparing, at the job scheduling unit, the resources required to execute the job and resources available based on the status information signals [in] (col. 3, lines 42-67; col. 4, lines 1-67; col. 5, lines 1-29).” Office Action at p. 7. Even assuming this assertion is true, which Applicants do not concede, Kubo fails to teach or suggest the claimed “scheduling ... based on ... the physical topology,” as recited by independent claim 10.

Because Fukuda and Kubo, taken alone or in combination, fail to teach or suggest each and every element recited by claim 10, no prima facie case of obviousness has been established with respect to this claim. Claims 11 and 12 depend from independent claim 10 and therefore include all of the elements recited therein. Accordingly, Applicants respectfully request that the Examiner reconsider and withdraw

the rejection of claims 10-12 under 35 U.S.C. § 103(a) as being unpatentable over Fukuda in view of Kubo.

III. Conclusion

In view of the foregoing amendments and remarks, Applicants respectfully request reconsideration and reexamination of this application and the timely allowance of the pending claims.

Applicants respectfully request that the Examiner respond to every point raised herein in accordance with M.P.E.P. § 707.07(f). Should the Examiner continue to dispute the patentability of the claims after consideration of this Reply, Applicants encourage the Examiner to contact Applicants' undersigned representative by telephone to discuss any remaining issues or to resolve any misunderstandings.

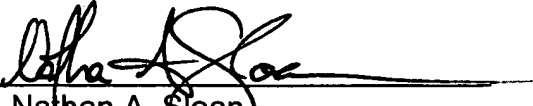
Please grant any extensions of time required to enter this response and charge any additional required fees to our deposit account 06-0916.

Respectfully submitted,

FINNEGAN, HENDERSON, FARABOW,
GARRETT & DUNNER, L.L.P.

Dated: June 7, 2006

By: _____


Nathan A. Sloan
Reg. No. 56,249
202.408.4312